A.2 Network Layer protocols:

A.2.1 X.25 over ISDN Basic Rate Interfaces (BRI) technology

ATIS committee T1S1 has completed its investigation into the Packet-Mode Communications issues identified by the FCC Report and Order 99-230[i]. The activity was addressed under T1S1's scope and charter to participate in the development of a joint T1-TIA standard on surveillance.

The T1S1 investigation specifically addresses the issues in FCC 99-230 and the request to identify capabilities which can be used to report Pen Register and Trap and Trace information for Packet-Mode Data Communications separately and distinctly from call or communication content. The focus of the investigation was on the X.25 over ISDN Basic Rate Interfaces (BRI) technology. The investigation was conducted on a technical merit basis and made no judgment with respect to legal issues regarding the applicability of the Communications Assistance for Law enforcement Act (CALEA).

A.2.1.1 Information that could be reported

A.2.1.1.1. Pen Register and Trap and Trace Information

For the X.25 Switched Virtual Circuits (SVCs) over ISDN Basic Rate Interfaces (BRI) technology, the packet handler will receive connection setup messaging from the subject. In these cases, the setup messaging exchanged between the subject and packet handler can be used as triggers for monitoring and reporting surveillance events. The following Pen Register and Trap and Trace information is available for X.25 SVC calls and could be provided to the Law Enforcement Agencies (LEAs) separately from the call content:

- Calling Party Number
- Called Party Number
- Answering Party Number
- Call Redirection/Call Deflection information
- Network User Identification (NUI)
- Recognized Private Operating Agency (RPOA)
- Called Line Address Modification Notification (CLAMN)

For the X.25 SVCs over ISDN Basic Rate Interfaces (BRI) technology, the existing J-STD-025 set of Call Data Channel (CDC) messages can be used to report surveillance events of packet-mode data communications during call setup, call progress, and call clearing (i.e., Origination, Termination Attempt, Redirection, Answer, Release). The J-STD-025 parameter set does not support the reporting of the Reverse Charging and Reason for Redirection information.

For X.25 PVC service, the ISDN packet handler uses a pre-provisioned connection across the packet network to deliver all transmitted packets between the subject and associate. In these cases, no connection establishment signaling is involved and, therefore, no end-

to-end routing information is available to the X.25 layer at the packet handler. The packet handler only maps the incoming ISDN connection from the subject to an X.25 PVC across the packet network. Therefore, Pen Register and Trap and Trace information for a PVC is not available at the X.25 layer.

For X.25 PVC service, only administrative records and operations personnel have knowledge of the end-to-end connection. This is because the operations personnel used the end-to-end information to provision the connection from the user to the packet handler (via an ISDN Permanent B-Channel Connection or D-Channel Connection) and then from the packet handler to an interoffice facility. At each switch in the PVC's path (which may cross network boundaries), the connection is mapped from one facility to another, until it is connected to the associate. Since the Pen Register and Trap and Trace information for X.25 PVCs cannot be derived by information at a given switch, the delivery of such information to LE over the J-STD-025 delivery interfaces cannot be automated.

A.2.1.1.2. Call Content

While only CDC messages should be sent for Pen Register type surveillances, Title III surveillances will require that call content be delivered over the packet Call Content Channel (CCC).

The call content information is available at the X.25 level. All X.25 packets should be intercepted and Pen Register and Trap and Trace information is not separated from call content before being replicated and transported over the packet CCC to LE.

The existing J-STD-025 set of messages to report the assignment and release of packet CCCs for the delivery of intercepted call content from packet-mode data communications can also be utilized (i.e., CCOpen, CCClose). The existing procedure in the J-STD-025 for reporting call content (Fast Select data) over the CDC can be used to report the transport of call content in X.25 SVC call setup, call progress, and call clearing packets.

The call content monitoring impacts for X.25 PVCs are similar to those described for X.25 SVCs. When monitoring call content for X.25 SVCs, the LEAs receive the necessary call parameters within the signaling packets that are also delivered over the packet CCC. However, for X.25 PVCs there are no signaling packets, and the call parameters are pre-provisioned for the connection. Consequently, when a court order requires call content monitoring for X.25 PVCs, the call parameters will need to be reported separately via a manual process, similar to the process for ascertaining the Pen Register and Trap and Trace information for X.25 PVCs. When call content is to be monitored on a X.25 PVC, certain call parameters may be needed to facilitate the LEAs processing of the call content data packets delivered over the packet CCC delivery interface. These call parameters include the packet modulo sequencing, the packet size, and the window size for the packet call. Without these parameters, it will be difficult or impossible for LEAs to properly extract the call content from the monitored packet-data communication.

A.2.1.2 Technical Impacts

The following Pen Register and Trap and Trace information is available for X.25 SVC calls but can not be reported to the LEAs because the existing J-STD-025 set of Call Data Channel (CDC) messages and parameters do not support the reporting of the information:

- Reverse Charging facility
- Reason for Redirection (as reported in the CLAMN facility)

When an X.25 PVC is re-provisioned to a different remote party (where the intercept subject is the local party), it may be problematic to ensure that LEAs receive timely notification. For X.25 PVCs that cross LATA (and usually state) boundaries, the ISDN service provider is not able to provide the identity of the remote party. Although the service provider may have noted the network address of the remote party in their records, this information has only nominal significance. Since the terminating party is in a different network, the local service provider cannot ensure that the X.25 PVC is actually connected to the network address listed in their records because network addresses do not have any real significance for X.25 PVC. The service provider can only confidently report the identification of the interoffice connection that is used to hand-off the X.25 PVC to the interLATA carrier. The LEAs would need to request that the interLATA carrier provide information about the connection; finally, the LEAs could then request that the remote ISDN service provider confirm the identity of the remote party.

Interception of packet services also does not guarantee that the packets have been received by the terminating system.

A.2.2 Asynchronous Transfer Mode

This section describes Asynchronous Transfer Mode (ATM) and its use in transporting voice telephony. It concentrates on the public network where ATM is predominately used as the bearer service for other, upper layer protocols that are concerned with the origination and routing of voice telephony calls. Therefore, the use of Switched Virtual ATM Connections is not described in this Appendix.

ATM is a switching method that uses fixed size units, called "cells," to transport information from the source to the destination. It is designed to be a general-purpose transfer mode for a wide range of services including, but not limited to, the transport of voice and data. ATM provides Layer 2 functionality in the Open System Interconnection (OSI) protocol layer model. Each ATM cell consists of a 5-octet header that defines the virtual circuit associated with the cell. Virtual circuits are defined by a combination of a Virtual Path Indicator (VPI) and a Virtual Channel Indicator (VCI). The remainder of the ATM cell consists of a 48-octet payload. In a typical public network, large numbers of virtual circuits are carried on the physical media.

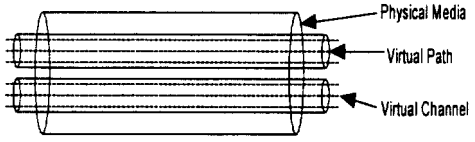

Included in the ATM header is a payload-type indicator that describes the cell as containing either user information or network management data. No information is included in the 5-octet header that defines the type of user data that is being transported.

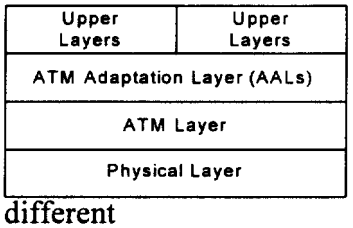

Separating the ATM layer from the user data is an adaptation layer that adapts the services provided by the ATM layer to those required by the higher layers. There are currently only three ATM Adaptation Layers (AALs) in common use and are described in this Appendix. Each different adaptation layer defines specific services to the upper layer applications that they are designed to transport.

Providing voice telephony over any bearer service requires the use of an Interworking Function to map the user's voice signals into whatever protocol the bearer requires.

Voice telephony in an ATM network can be provided by AAL1 (Circuit Emulation Service), AAL2 (Voice Over ATM) and AAL5 (Variable Bit Rate).

ATM adaptation Layer AAL1 defines how Time Division Multiplexing (TDM) type circuits can be emulated over an ATM network. AAL1 supports emulation of DS1, DS3, and nxDS0 circuits and is used primarily to provide transport of PBX trunks.

ATM adaptation Layer AAL2 defines a method to provide variable length packet payload multiplexing within a single virtual circuit. One of the major advantages to using AAL2 is that it provides for multiplexing of voice packets into a single ATM cell. Either AAL2

or ATM adaptation Layer AAL5 is used to transport signaling data (DTMF, Common Channel Signaling, etc.) to the Voice Gateway.

ATM adaptation layer AAL5 defines a method to provide variable bit rate services primarily for data applications where the bursty nature of the applications can tolerate variations in delay. Voice over Packet (VoP) may be carried over AAL5, but the information above the ATM Layer appears as data to the ATM Layer.

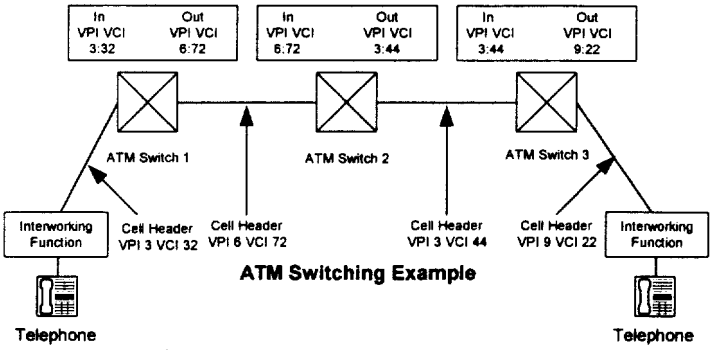

As ATM cells arrive at the ingress to the network, each cell header is examined to determine if the cell contains network management or user information. If the cell contains user information, the VPI and VCI are looked up in a routing table to determine the outgoing facility to use in transporting the cell. Because the outgoing facility may use a different VPI/VCI combination, the node must replace the VPI/VCI with the new values. Cells continue in this way, moving from node to node until they reach their final destination. At no time during this process of *relaying* cells does the network know anything about the upper layer protocols that are being transported. Note that the VPI/VCI only indicates the association between the adjacent nodes and not the end-to-end connection. Nothing about the VPI/VCI defines the final destination of the cell or the user data that is being transported. Thus, Pen Register and Trap and Trace information is not available. ATM switching nodes are designed to relay the ATM cells as quickly as possible. This design effectively prohibits the ATM switch from reassembling and examining the user information encapsulated in each cell due to the processing demand and implementation that would be required. Within an ATM network, it is not technically feasible to extract the upper layer protocol, which contains the information of interest to the LEA from the ATM cell stream. It may be technically feasible, however, to extract call content information at the ingress and egress service interfaces.

A.2.3 Internet Protocol

The IPv6 RFC 2460 is available and may become widely deployed in the future. However, as its deployment is limited at this time the JEM considered only IPv4 impact on CALEA and did not consider the impact of IPv6.

A.2.3.1 Introduction

This section analyzes the areas identified in the main text of the section as they apply to networks using the Internet Protocol.

This section first discusses architectural principles of the Internet and the Internet Protocol that apply to the analysis of CALEA. It then follows the basic organization of the main text in that it investigates what information can be delivered and the technical issues of delivering that information both for a provider that supports a Call Management System and a provider that only supports IP transport.

Although the Internet (and IP) supports many applications other than Call Management Systems, due to the special consideration given to voice applications in the JEM this section only deals with Call Management Systems.

A.2.3.2 Scope

This section concentrates on the Internet Protocol and related protocols.

While the Internet Protocol can be carried by a multitude of underlying protocol technologies (e.g., leased line, dial-up modem, ATM, Frame Relay, X.25, etc.), this section does not consider the implications of CALEA on the underlying protocols. This is left to the appendix for the specific technology.

In addition, this section does not provide special considerations to any applications other than Call Management Systems that run over IP.

A.2.3.3 Architectural principles related to CALEA

This section briefly describes architectural principles of the Internet that apply to the CALEA analysis.

There are plenty of tutorials and books on IP and routing available. It is assumed the reader is familiar with the operation of the Internet and the suite of Internet protocols. However, discussions of general principles that affect the issue at hand are included.

A.2.3.3.1. End-to-end principle

"The network's job is to transmit datagrams as efficiently and flexibly as possible. Everything else should be done at the fringes." [RFC1958]

The end-to-end principle is simple, but powerful. It recognizes that there are many functions that only make sense to implement in hosts at the edge of the network. Examples are reliability, security (encryption), etc. This is in marked contrast to other protocols and networks that have been developed over the last 100 years which attempt to subsume these functions into the network.

As an example, the Internet assumes that, in general, reliable data transfer is assured by the end systems instead of the network. What this means is that any retransmission due to packet loss is done end-to-end instead of inside the network. As a counterexample, X.25 and similar protocols provide retransmission on a hop-by-hop basis.

Quoting from [Saltzer], "The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the endpoints of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete

version of the function provided by the communication system may be useful as a performance enhancement.)"

This does not mean that a service provider cannot offer enhanced services over such a network. What it means is that the enhanced services are either provided as part of the transport mechanism (e.g., Quality of Service) or are enabled on hosts reachable via the Internet. Examples of such systems are DNS, email servers, Web servers, etc. The customer reaches these services just like it would any other service not offered by the service provider (i.e., via IP).

A by-product of this principle is that the IP transport network will not necessarily know what applications are being run over the network since there is no "setup" in which the network participates. In fact, the network is designed not to know what application is being carried. The applications run end-to-end and the network just routes packets. Where the network does provide an application-level service, it is through a host that it manages that communicates end-to-end with the customer's host over the provider's network.

The end result of this is an explosion of innovation in applications. Anyone with a link to the Internet can develop and offer new and innovative services to anyone else on the Internet (e.g., Napster). This applies to voice applications just as much as any other applications.

A.2.3.4 Security

Security encompasses many areas including, but not limited to, encryption. In general, the end-to-end principle applies to security (i.e., end-systems are responsible for their own security). However, service providers can provide assistance in some areas of security. For example, even though two end-systems may have adequate security in their locations and use strong cryptography between them, unless they have the cooperation of their service provider they are susceptible to denial of service attacks from third parties that flood their link to the Internet such that communications is degraded.

As discussed in Section 6 of [RFC2804], the introduction of capabilities for electronic surveillance tends toward making the network itself less secure, even when the capability is not being exercised. Much effort is underway in the industry to make the Internet more secure, not less. Development of specific protocols and methods for delivering an end-user's information to a third party without the knowledge of the end-user does not contribute to making the Internet more secure.

Encryption has been a controversial topic for a long time. However, with the advent of more powerful computing devices and more powerful and available encryption algorithms, strong encryption is now technically available (but possibly not legally or politically available) to most people on the Internet. Following the end-to-end principle, encryption should take place between two hosts, not in the network itself. The network may use encryption for its own purposes, but the hosts using the network ultimately must take care of themselves. The end systems may have a trust relationship with the service provider that enables the service provider to share in the security mechanism and encryption of data; however, the fact that CALEA may require the service provider to provide decrypted information (or keys) to an outside body (i.e., LEA) without the user's

knowledge will force the service provider to implement mechanism that could make the system less secure than it might otherwise be, even if the subject is not under surveillance.

A.2.3.5 Encapsulation

"IP on everything" - Source Unknown

The Internet Protocol is designed to operate over a wide variety of network technologies and protocols. In fact, the term internetwork (thus internet) derives from the fact that it was designed to interwork over multiple networking. Normally, an encapsulation method is defined for how to run IP over a particular networking technology. Encapsulation methods have been defined by the IETF for a wide variety of networking technologies (e.g., HIPPI, X.25, Frame Relay, ATM, FDDI, Ethernet, Token Ring, Arcnet, leased line, dialup, etc.). As IP is routed from one network type (e.g., SONET) to another (e.g., Ethernet) it is encapsulated into a different mechanism that is usually specific to a particular network type.

IP also encapsulates upper layer protocols inside its data field. These upper layer protocols are usually transparent to the IP layer and to the Internet between two hosts. IP provides a Protocol ID that identifies the protocol contained in its data field. TCP and UDP are two predominant protocols that run over IP; however, other protocols can be run directly over IP (e.g., IPSEC, IP, RSVP, SCTP, etc.). Over 100 Protocol numbers have been assigned by IANA for use in the IP Protocol ID field. Normally for user data transfer, the Protocol ID does not identify the application the hosts are running. The applications normally run over a transport protocol (e.g., UDP or TCP) that runs on IP. The end systems can identify which application a particular packet is destined for by the TCP (or UDP) port number. There are several thousand port numbers currently registered with IANA for use with TCP or UDP.

Thus the application data is usually encapsulated in UDP or TCP, which is encapsulated in IP, which is encapsulated in a link or network specific mechanism.

A.2.3.6 Connectionless Orientation

The Internet Protocol is a connectionless protocol. In general, each packet contains all the information needed to route the packet from one host on a network to another host on the same or different network. Each packet is routed through the network(s) independent of the previous packet and may take a different path through network(s) than a previous or subsequent packet. There is no explicit setup mechanism between a host and the network to provide communication between two hosts. There is no "call" in IP. Loop Start in the analog telephony world and Q.931 in the ISDN world are examples of signaling protocols between a host (e.g., telephone) and network that set up connections (i.e., calls) between two hosts (e.g., two telephones).

As mentioned earlier, a service provider can provide services by deploying hosts in the network to which customers' hosts can communicate. A customer's host can request service from the provider's host; however, the network provides the connectivity for the

packets. An example of this is DNS. A host wants to communicate with another host, but only knows the host name and not the IP address. The host sends a query to the DNS server via IP and the DNS server responds with the destination host's IP address. This is still a connectionless service.

A.2.3.7 Boundaryless

"There's a freedom about the Internet: As long as we accept the rules of sending packets around, we can send packets containing anything to anywhere." [Berners-Lee]

One by-product of the above principles is that IP inherently has no concept of geopolitical boundaries. While a network's design may provide some loose constraints as to what path a packet may take, there is usually no guarantee a packet will take a particular path at any particular time.

For example, a host in Tuscaloosa, Alabama may download a file from another host in Mobile, Alabama. There is no guarantee that the packets in this download will stay in the state of Alabama. They may transit part of the network in Mississippi, Florida or any nearby state. This is considered part of normal operation of the Internet. Therefore, information that may be reasonably available in a connection-oriented network may not be available in an IP network.

A.2.3.8 Call Management System

The main text contains general information for packet-mode technologies concerning information derived from a call management system. This section contains information specific to the Internet Protocol.

Call Management Systems don't exist for the Internet Protocol. However, call management systems exist for applications that run, end-to-end, over IP. In general on an IP network, a call management system is a host attached to the IP network running call management protocols end-to-end over IP to its clients. For VoIP applications, the encoded voice stream is also carried over UDP/IP. In VoIP, the IP packets carrying voice are usually carried directly between the two endpoints involved in the call. The Call Management System is not involved in transporting the voice packets.

A.2.3.8.1. Information that could be reported

The information discussed in the main text also applies to IP-based Call Management Services. Since the information available from Call Management Services tends to be specific to the application (i.e., Voice) as opposed to IP, the information itself is pretty much the same as the main text. The call events described in J-STD-025 are examples of information that can be reported.

In a network that provides a call management system, only call events that are triggered by messages between the target and CMS are available.

For IP, the call content flow within the immediate network can be characterized by the source and destination IP address and source and destination UDP port numbers negotiated between the CMS and target during call establishment. This only applies to

call control protocols that exchange IP address and port numbers with the CMS. This applies to most VoIP signaling protocols currently defined, but may not apply to future applications.

The following are IP-specific items:

IP information for call management protocol:

IP address used by target
TCP or UDP port number used by target
IP address used by Call Management System
TCP or UDP port number used by Call Management.

IP information for voice packet stream

IP address used by target for voice packet stream
TCP or UDP port number used by target for voice packet stream
IP address used by target's associate for voice packet stream
TCP or UDP port number used by target's associate for voice packet stream

Other call events similar to those defined in J-STD-025 may be available to the extent that the Call Management supports similar services (e.g., Call Forward).

A.2.3.8.2. Technical issues

The issues addressed in the main text also apply to IP. This section discusses issues specific to IP.

A.2.3.8.3. Location and Ownership of Call Management System

Since it runs over IP the Call Management System can be located anywhere on a global IP network. It can be provided by the Service Provider that also provides the Internet access service to a customer, by another Service Provider on the Internet or by a customer on another Internet Service Provider. The CMS may not be within the same jurisdictional boundaries as the client host to which it is providing service even if the two endpoints involved in a call are within the same jurisdictional boundaries. The only real requirement is that the users have reachability via IP to the CMS. In the extreme case, the CMS could be in a different country from the hosts to which it is providing service.

This follows the principle laid out in Section A.2.2.7

Given that Call Management protocols are end-to-end protocols over the Internet, a Service Provider will only have access to call events detected on its Call Management System. This section only discusses issues with what information can be gleaned from a Call Management System operated by the Service Provider.

The Call Management Server in an IP network can only report events based on packets that are terminated on or originated from the Call Management Server. If the target knows the destination address of the person it wants to call or if the target uses a Call Management System not under the control of the service provider, the target can establish a VoIP call without the knowledge of the service provider's Call Management System. Call Events for such calls will not be available to the service provider.

Even when the target uses the provider's Call Management Server, not all call events for call manipulation may go through the Call Management Server. For example in the middle of a call the endpoint under surveillance may exchange call control information with a Call Management Server not under the control of the service provider. These packets will not necessarily go through the Call Management Server, but will be routed normally as data packets.

A.2.3.8.4. Call Management Protocols

The Internet places no restriction on the protocols used between the CMS and client for managing calls. In today's PSTN, the base protocols used for call control are limited to a small number due to the technology and the small number of providers. Each country or network may define its own variant of the base PSTN protocol but they all have the base in common.

The following is a list of some of the protocols defined for call management over IP by various industry groups

- H.323,
- SIP,
- H.248/megaco,
- MGCP,
- PINT (based on SIP)

More protocols are being developed. Each of the above listed protocols is fairly flexible in allowing different services based on the core protocol. Therefore, nailing down a complete, fixed set of call events that are available via each protocol is close to impossible.

In addition, some of these protocols (e.g., SIP) can be used for provision of information services. In fact, the same CMS host could offer information services using the same protocol at the same time as offering VoIP service.

Since the CMS and the client usually run on open computing systems, new call control protocols are usually easily downloaded and installed. The CMS and client don't have to run standard protocols as long as they agree with each other what protocol to use.

The ubiquitous HTTP (i.e., protocol used to support the World Wide Web) is also being used by various entities to offer VoIP services such as "click-to-dial". In this case, there is not necessarily a specific call control protocol and the CMS is a web server.

In some protocols (e.g., SIP, H.323), the information exchanged between the client and the CMS may only be sufficient to resolve an identifier such as an email address to an IP address which the client uses to negotiate the call further. In this case, subsequent call events may not be available to the CMS.

Although the call control protocols have UDP/TCP port numbers assigned to them via IANA, there is no hard requirement in the Internet to use these port numbers. The port numbers used for call control is a bilateral agreement between the CMS and client. Most

of the time, applications will utilize the IANA-assigned port numbers at least for the initial information exchange; however, this is easily changed by agreement.

The port number negotiated between the CMS and the end-systems for the actual voice stream is variable and dynamic. It is assigned to each end of the call on a call-by-call basis. In addition, some call control protocols may use a dynamically assigned port number for negotiating supplementary and other services.

In conclusion:

- There is not a clear, fixed definition of CMS for VoIP.
- There is not a complete, fixed, limited set of services defined for VoIP. However, a limited, fixed set of services can be defined that might be available by most providers (e.g., connect, disconnect, forward, transfer, etc.).
- There are multiple protocols a client can use to communicate to the CMS.
- The UDP or TCP port numbers used between CMS and client for call control are via bilateral agreement. Most of the time they are IANA-assigned.
- The CMS used for VoIP support can also offer Information Services at the same time using the same protocol.

A.2.3.8.5. Service Paradigms

The CMS does not necessarily follow any traditional paradigm in terms of the services it offers. Some Service Providers will use a CMS to offer VoIP services that mimic today's PSTN as closely as possible, thus recreating the current telephone system on IP. In this case, it is reasonable to expect that call events similar to those defined in J-STD-025 or PacketCable(TM) may be available. Other Service Providers, or end-users on the Internet, may provide innovative services that are not possible or available on today's phone system and may not offer many of the services that are available on today's phone networks. In fact, some of the new services may not be recognizable as traditional voice services and could be interpreted as information services. This is the anticipated result of the end-to-end principle, which enables anyone attached to the Internet to develop and offer services.

The distinction between an "electronic messaging service" which as defined by CALEA includes audio (e.g., voice) and is not included in CALEA requirements and a VoIP service will tend to blur as time goes on and the market develops. For example, an end-user could send an email with an audio attachment. The receiver of the email could listen to the attachment and send back an email with an audio attachment as a response. This may fall under the "electronic messaging service" and could also be seen as a voice conversation with a long delay.

The amount of control a Call Management Server exercises over a target's endpoint varies greatly depending on the standard used and how it is used. It can exercise very little control, such as in some SIP or H.323 RAS cases, or detailed control such as in H.248 or MGCP. Thus, the amount of information available will depend on the service offered by the service provider and by the protocols used for providing the service.

If the Service Provider is providing a traditional telephony service over IP, then call events such as those defined in J-STD-025 or the PacketCable specification should be available.

A.2.3.8.6. Redirection of Calls

As discussed previously, packets carrying the call content are not guaranteed to follow a particular path through the Internet. Nor are they guaranteed to stay on the provider's network even if the two endpoints are on the provider's network. If the target is communicating with another endpoint off the provider's network and redirects the call to another endpoint off the provider's network, then most likely all packets involved in the redirected call will not transit the provider's network (although it is possible they will). Therefore, these packets are not available to the original service provider and cannot be provided. In addition, since the call is redirected, the service provider's Call Management Server is no longer involved and may not have access to any call events related to the redirected call.

A.2.3.8.7. Target Identification

In order to provide the information required, the correct target must be identified.

If the service provider has a relationship with the target and the target uses the service provider's CMS, then it will probably be able to identify the target via either a login, pre-assigned IP address or other pre-assigned identifier (e.g., calling party number).

If the IP address of the target is known, it can be used to identify call events from the target. Usage of IP addresses has the same issues discussed below for IP Transport.

However, some services offered on the Internet may not require a specific relationship between the provider and the client. For example, these may be ad-hoc services offered via the World Wide Web. The provider of the service may not know who is using the service, much like a web server today. In this case, the only information it may have about the target is the information the target sends it. The target's IP address would be one constant that could be used to identify the target. Again, correlating the IP address to the target will have the same difficulties as defined below for IP transport.

In traditional PSTN telephony, the telephone number can be used to identify the called and calling party. It is recognized that this usually identifies a physical telephone that anyone can use, including people that are not under surveillance. In VoIP, a telephone number may not be used to identify the caller. For example, SIP allows the use of email-like addresses (e.g., foo@bar.com) to identify the called and calling party. Therefore, existence of telephone numbers is not guaranteed on the CMS.

A.2.3.8.8. Performance and Complexity

The performance impact on the Call Management System will depend on the services offered and the information required. If information similar to J-STD-025 is provided, the performance implications will be similar.

If the protocol used to deliver the information to the LEA is significantly different from the protocol used by the CMS for its other communication, it could have an affect on the complexity and performance of the system. For example, if the CMS is using IP to

provide service but must use X.25 to communicate to the LEA, the separate drivers and software to run X.25 may introduce additional complexity. In addition, if the data formats used to deliver to the LEA are substantially different, then this will also introduce additional complexity (e.g., a CMS that does not use ASN.1 to provide service but must use ASN.1 to provide CALEA information).

A.2.3.8.9. Delivery Format

The various call control protocols listed above utilize different encoding mechanisms. For example, SIP is text encoded with syntax defined in Augmented Backus-Naur Format (ABNF) while H.323 is binary-encoded with syntax defined in ASN.1. H.248 allows both. Thus, it is not possible to define a delivery protocol that is consistent with the encoding mechanism of each protocol.

However, it is possible to define a delivery mechanism for all the identified protocols that run over IP and reduce the complexity of using multiple network protocols for delivery.

A.2.3.8.10. Points of Intercept

For CMS, the point of intercept would most likely be on or near the CMS.

A.2.3.9 IP Transport

This section focuses on information available from a service provider that is offering IP transport. In this case, the subject is not using a Call Management System offered by the service provider.

A.2.3.9.1. Information that can be reported

The IP transport provider utilizes the information in the IP packet for providing service. In general, routers are optimized to operate on the IP header for providing service. Layer 2 switches are optimized to operate on Layer 2 headers to support transmission of IP.

A copy of the IPv4 header is shown below:

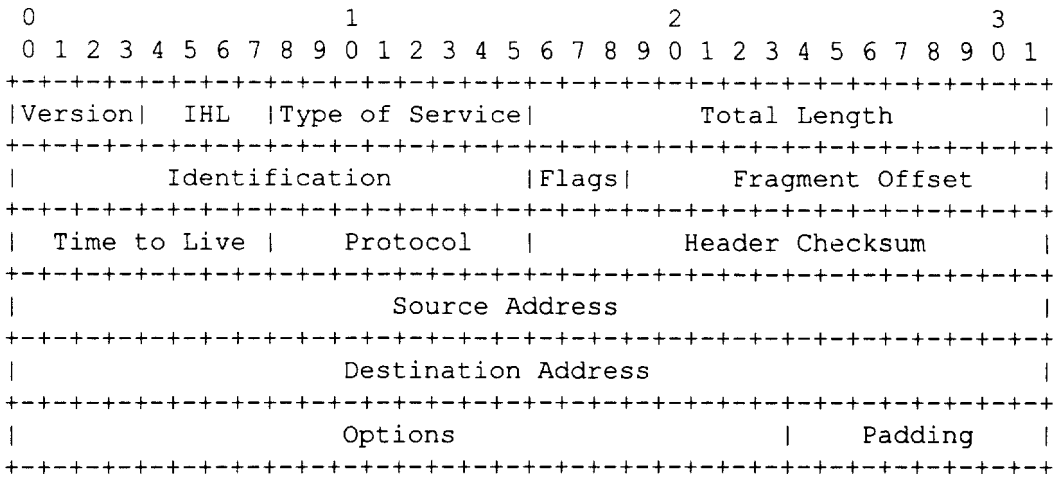

The information in this packet header that identifies the source and destination endpoints is the Source and Destination Address fields.

Since the provider in this case does not participate in any call control, it does not have access to information identifying a call. All it does is transport packets. Note that this discussion also applies to a Service Provider that provides a CMS when the target is not using the CMS.

The next layer protocol is identified by the Protocol field. The valid values for the Protocol field are defined by IANA. The next layer protocol could be UDP, TCP or could also be IP. There are several hundred Protocol Numbers defined by IANA for protocols running over IP.

The following three cases were identified at JEM I concerning information that can be provided:

1) Transmission of entire packet stream: In this case, the provider transmits the entire packet stream to and from the target to an LEA and the LEA uses minimization to extract the information to which they are legally entitled. However, it should be noted that this is similar to the original solution, which caused the FCC to request this report. In other words, there is no separation of Pen Register and Trap and Trace information from call content. In this section, transmission of the entire packet stream means transmission of the packets whose source or destination IP address match the target's.
2) IP source and destination address: This is the source and destination address contained in the IP header. In the case of tunneling, this would be the information contained in the outermost IP header. This could also include the Protocol ID field.
3) Extraction of information in the IP data field: This consists of the provider examining the data field of each IP packet in the packet stream to discover "Pen Register and Trap and Trace information." Since IP uses successive encapsulations to carry data, the question of how deep into a packet a Service Provider must go to retrieve the information. Examples of information discussed included:
   - TCP or UDP port numbers; and
   - Call Control (e.g., SIP Invite, H.323 SETUP) information extracted from data field.

A.2.3.9.2. Technical Issues

This section discusses the issues surrounding delivering the information identified above as well as several general issues related to IP technologies.

A.2.3.9.3. IP Packet Fragmentation

The maximum packet sizes supported by the various networks in the path from source to destination may vary. When a router receives a packet on one interface that is larger than the maximum packet size of the interface to which the packet is routed, the router might fragment the large packet into multiple smaller packets. A host might also perform this fragmentation, but it is recommended that the application not generate packets that are too big.

If an IP packet is fragmented, only the first packet will contain the upper-layer protocol headers. For example, only the first packet will contain the TCP header and thus the port numbers. The subsequent packets will not contain the port numbers. Therefore, port numbers may not be available in each IP packet.

A.2.3.9.4. Target Identification

This is a general issue related to all three cases identified in JEM I in addition to the Call Management System case. In any of the above cases, the LEA must identify the target.

There are two general issues involving target identification:

Use of Network Address Translation; and

Use of Dynamic IP Address Assignment.

A.2.3.9.4.1 Use of Network Address Translation

The global Internet uses globally unique IP addresses. However, provision is made [RFC1918] for usage of private addresses. Two non-overlapping private networks may use the same private IP address space.

When end-systems on different networks using private addresses (or an end system on a private network and an end system on a public network) need to communicate, a method called Network Address Translation is used to translate the IP addresses between the two networks. Therefore, the IP address contained in an IP packet in one network may be different from the IP address for the same packet in another network. The address translation between two networks can be fixed (address in one network statically mapped to an address in another network) or dynamic (address in one network dynamically mapped to one of a pool of addresses in the other network). In the case of dynamic mapping, the address seen in one network may be mapped to a different address in the other network depending on time of observation.

A.2.3.9.4.2 Use of Dynamic IP Addresses

IP addresses can be "leased" for a period of time, released and reused. These addresses are usually called "dynamic" IP addresses as opposed to "static" IP addresses, which are assigned to a user for an extended period of time. A dynamic IP address only identifies an endpoint for the duration of usage by that endpoint. Several mechanisms have been defined to assign IP addresses dynamically. The most common are Dynamic Host Control Protocol (e.g., in LANs and cable networks) [RFC1541] and Internet Protocol Control Protocol (A control protocol used by PPP endpoints, e.g., ., in dialup or xDSL networks) [RFC1332]. Therefore, an IP address may only identify the target or its associate during the duration of a particular connection or session. The target and its associate may have a different IP address at a different time, even if connecting from the same location.

In some of the dynamic addressing schemes, the user authenticates with the network before being assigned an IP address. For example, in dialup Internet access the user usually has to authenticate using PPP before it is assigned an IP address. The user database is stored in a server and the Remote Authentication Dial In User Service (RADIUS) [RFC2138] protocol is used between the network access device and the server to determine if the user is allowed to access the network. However, in many cases the IP address is allocated from a pool on the network access device and not from the RADIUS server. Therefore, even the RADIUS server that authenticates the user may not be aware of address assigned to the user. If the provider uses RADIUS Accounting [RFC2139], then the accounting messages sent from the network access device to the accounting server may contain the address information. RADIUS is an example of an authentication and accounting protocol and is not necessarily the only one used in a network.

A.2.3.10 Transmission of entire packet stream

In the case (case 1) of transmission of the entire packet stream, there is no separation of Pen Register and Trap and Trace information from content since the provider is providing all the data to and from the target.

In order to transmit the entire packet stream to the LEA, the packet stream must be extracted from the aggregate packet stream. There are a several methods for doing this:

- Packet replication: In this method the service-providing equipment (e.g., router) replicates the packet stream to/from the target and transmits the packet stream to the LEA (possibly from a different interface). A simple example of this is to connect to a span port on an ethernet switch.
- "Sniffing" the packet stream: In this method, the provider or LEA connects equipment to the physical medium to extract the packet stream from the aggregate packet flow. A simple example of this is to tap into a coax ethernet cable and copy off packets.

If packet replication is performed on the service-provider equipment, the capacity used for replication will not be available for providing service. This may affect service to the provider's customers including the target. Packet replication may require hardware support that is not available in all equipment.

"Sniffing" the packet stream requires non-obtrusive physical access to the transport medium. Some physical media (e.g., fiber optic) is not conducive to a non-invasive tap unless a tap is preinstalled for such activity.

There are several issues with transporting the entire packet stream. One is based on the service provided. For example, a virtual private dial service encapsulates the PPP packet from the customer into an IP packet (using L2TP) destined for a gateway into another network. In this case, the Network Access Equipment does not normally assign or look at the IP addresses in the customer's packets. In addition, multiple user sessions can be multiplexed into one tunnel to the remote side. Monitoring behind the NAS would require specialized equipment to de-encapsulate the tunneled packet (and possibly de-encrypt) to extract the original IP addresses.

A.2.3.11 Transmission of IP Source and Destination Address

To transmit the IP source and destination addresses (case 2) requires the equipment to read the IP header, extract the source & destination address and deliver that information to the LEA.

Service providing equipment may not be designed to extract IP header information and deliver it to the LEA. Other equipment may be able to do this. In any case, the processing capacity used for delivering the header information is not available for routing packets. The amount of load on the system will depend on the system. There is a multitude of vendor equipment of different types deployed for Internet service and each one would have to be tested to determine what load it would support.

On the other hand, specialized equipment exists today that can extract the IP header information and some other information (e.g., TCP or UDP port number) from a real-time stream.

A.2.3.12 Extraction of Information from Packet Stream

The extraction of information from a packet stream for delivery to the LEA was one option (case 3) discussed at the JEM and other fora. For example, the provider would be required to monitor the packet stream, detect a call control packet containing Pen Register and Trap and Trace information (e.g., for VoIP), extract the Pen Register and Trap and Trace information from the packet and deliver it to the LEA.

Routers supporting service on the Internet typically only make routing decisions based on the IP addressing information. Service providing equipment is not generally designed to look past the IP headers (some may look at TCP or UDP port numbers for filtering) when switching or routing packets. Any processing capacity used for extracting information from a packet stream is not available for routing packets. Given the increase in capacity of Internet connections and that systems generally run at peak load much of the time, there is very little capacity to monitor data fields.

An alternative is to use equipment that is not providing service to the customer but has access to the data stream (e.g., via a port on an ethernet switch). This equipment could acquire the information required and deliver it. This would require extra equipment by the service provider and new operating procedures. In addition, any time new equipment

is added to a network it introduces the possibility of errors and misconfiguration and can disturb the functioning of the network.

However, as noted in the main text, there is no reliable method for determining the Pen Register and Trap and Trace information when monitoring a packet stream. If given a specific IP address and port number and if encryption or tunneling isn't used and if the call control protocol is identified then it might be possible to extract the call control information for VoIP calls. However, there is no guarantee that the session is a telecommunications service or an information service. It would be similar to requiring telecommunications carriers to monitor inband communications, detect and demodulate modem tones, detect and decode the information carried in the modem signal and extract Pen Register and Trap and Trace information that may be carried. This is not equivalent to detecting in-band DTMF.

A.2.3.13 Tunneling

One of the issues with identifying the target and providing information is the use of tunneling. In essence, tunneling is the act of encapsulating network protocol packets into IP packets to be routed across the network. In this section, the tunneling of IP packets is considered. The tunneling method defines a mechanism to encapsulate IP packets inside other IP packets. The outer IP header is used to route the packet across the network. The source and destination addresses of the outer IP header identify the tunnel endpoints. The IP addresses in the IP header may not be the IP address of the final endpoints. For example, the tunnel endpoint could de-encapsulate the IP packet and route it onward using the encapsulated IP address information.

There are several methods that can be used for tunneling IP packets across an IP network: Generic Routing Encapsulation [RFC1702], IP-in-IP [RFC1853], Layer 2 Tunneling Protocol [RFC2661], IP Encapsulating Security Payload (ESP) [RFC2406], etc.

The network routes packets based on the outer IP headers and not on the inner headers. In some cases, such as in IPSEC ESP [RFC2406], the encapsulated IP packet is encrypted and isn't available even if the service provider could sniff into the data packet.

For tunnels originated from the target, the destination IP address in the IP packets from the target and the source IP address of IP packets to the target are not necessarily the IP address of its associate. This IP address could be a tunnel endpoint, which will de-encapsulate and route the tunneled packet onward. The source IP address of IP packets from the target and destination IP address of IP packets to the target will normally be the IP address of the target in order for packets to be routed to it properly.

For tunnels originated in the network, the original IP headers may be available. However, in some cases, such as L2TP, the original IP addresses may not be readily available. In L2TP, the network access device encapsulates all PPP frames from the user into IP packets into another IP packet destined for another location. In this case, the IP addresses of the outer IP header will be the IP address of the network access device and the remote tunnel endpoint. The target's IP address will not be in the outer IP header at all. The IP address of the target is assigned by and is only seen by the remote tunnel endpoint. The tunnel from the network access device to the remote tunnel endpoint may be shared by many users including the target.

A.2.3.14 IP Address Spoofing

Another issue that will affect the integrity of the information provided to the LEA is IP address spoofing.

A destination IP address must be authentic or else it cannot be routed to an endpoint. However, a source IP address may or may not be authentic since it is not required for the network to route packets to the destination properly. A valid source IP address is required for the destination to transmit packets properly back to the source. However, there may be cases in which an entity may not care about receiving responses.

Although there are several methods available to do so, in general network access devices do not check for invalid source IP addresses before accepting packets into the network and forwarding them. Therefore, this allows endpoints to spoof the source IP addresses.

There are two cases when dealing with spoofed IP addresses and CALEA:

Target spoofing source IP addresses to a destination

Associate spoofing source IP addresses to the target

In the first case, if the provider is keying on the source IP address in order to provide information to the LEA and the target is using a shared line to access the network, the provider would not necessarily detect and supply information on IP packets with spoofed source IP addresses from the target. A consequence of this is that the target could carry out a Denial of Service attack on a remote host without it being detected by the tap.

In the second case, a remote endpoint could send packets to the target with spoofed source IP addresses. This could result in several things happening:

Under a Trap & Trace order, the remote endpoint could cause the LEA to investigate users who have no relationship with the target by spoofing their IP addresses.

Under Title III, the remote endpoint could incriminate the target (and other users) by sending illegal material to the target with spoofed source IP address.

Although the information required to carry out the above attack is not necessarily readily available (e.g., the fact the target is under surveillance, the IP address of other endpoints, etc.), it is possible for a sophisticated user with knowledge.

Software to spoof IP packets is readily available on the Internet. Tracing a packet stream with spoofed source IP address back to the originator is extremely difficult to do on the Internet, especially if the packet flow is intermittent.

A.2.3.15 Delivery format

In defining the delivery format for IP, the characteristics of internet connections should be taken into account.

The connection speeds from providers to their customers is continuing to increase. The delivery vehicle for the information collected may have to be substantially different for a customer connected via Gigabit Ethernet to one connected via a dial-in modem.

For delivering IP addresses, the provider could provide large amounts of effectively the same information for large data flows since the IP addresses don't usually change for a data flow. Alternatively, the provider could supply the information once per data flow. The latter option could reduce the amount of information transmitted to the LEA and the strain on the system.

A.2.3.15.1. Transport Protocol

In defining the protocol used between the POI and the LEA, the following characteristics of the transport protocol must be taken into account:

- Reliability: While TCP will provide a reliable connection between the LEA and the POI, it will also require more processing in the end-systems. If a real-time packet stream is required to the LEA, TCP may not be suitable since it is not designed for real-time transport. A UDP based protocol will require the least processing; however, UDP does not provide reliability.
- Security: If IPSEC is used between the POI and the LEA, the processing requirements of IPSEC must be taken into account.

A.2.3.16 Performance and Complexity

Due to the different types of network access, the different types of equipment and different vendors, acquiring hard performance numbers on the impact of providing information is not possible. It would require testing of each piece of equipment in each scenario.

However, some general principles can be considered.

Introducing new features and new code into a system introduces complexity into the system along with probable errors. For service providing equipment, since monitoring streams is a real-time activity that touches the mainline of packet forwarding, the complexity is added directly to the part of the system that can least afford added complexity. If the monitoring is provided in off-line equipment that is monitoring the communications then the system can be optimized for monitoring and filtering functions without necessarily affecting service to customers.

A.2.3.17 Points of Intercept

The point of intercept for an IP transport provider would preferably be close to the edge of the network. The closer to the core of the network, the heavier the traffic flows and the more random the traffic patterns.

The point of intercept should be flexible based on the provider's network design and equipment capabilities. The following are possibilities:

One location for the Point of Intercept is the first equipment that routes the target's traffic. Another location is an aggregation router through which the target's traffic flows. However, as discussed previously, performing packet monitoring an delivery from equipment that is providing service to a customer detracts from the resources available to service customers.

Another Point of Intercept is within the access POP for the provider. The raw packet stream can be provided via a port on a POP switch (e.g., ethernet switch)

or via a line monitor. Non-service providing equipment can be used to take the raw packet stream as input and provide the required CALEA information as output.

No matter where the Point of Intercept is located, some correlation has to be provided to correlate the target's current IP address with the surveillance stream.

A.2.4 Frame Relay

This section describes Frame Relay (FR) technology and its use in transporting voice telephony. It concentrates on the public network where FR is predominately used to transport packetized data. Traditionally FR has been used to transport LAN to LAN and legacy traffic such as bisync and SNA. Recently, non-traditional uses (Voice over Frame Relay [VoFR]) have begun to materialize.

Early transport of packetized data made use of X.25. X.25 includes considerable overhead that Frame Relay was designed to overcome. The major differences between X.25 and Frame Relay are:

- Logical connection multiplexing and switching takes place at Layer 2 rather than Layer 3;
- Frame Relay leaves end-to-end flow and error control to upper layers; and
- User data and call control signaling are carried on separate connections.

The advantage of Frame Relay lays in the fact that the communication process has been streamlined to take advantage of modern transmission systems that are less prone to error. By lowering the overhead necessary to transport data, increases in throughput and decreases in delay have been realized.

Frame Relay switching is best understood by examining the frame format:

Flag 1 octet	Address 2-4 octets	User Data Variable	FCS 2 octets	Flag 1 octet

The address field has a default length of 2 octets, but can be extended to 3 or 4 octets. Included in the address filed is the data link connection identifier (DLCI) that is used to define the virtual circuit number. This arrangement allows for the multiplexing of many virtual circuits onto a single physical media.

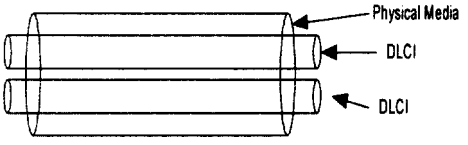

Voice over Frame Relay (VoFR) offers the promise of consolidating data and voice traffic over the same Frame Relay network. Recent advances in Frame Relay features, specifically to support real-time services such as voice, have made Frame Relay an attractive alternative for carrying voice traffic. The Frame Relay Forum introduced the Voice Over Frame Relay Implementation Agreement (FRF.11) in December 1998.

The Voice Over Frame Relay Implementation Agreement (FRF.11) includes features to support real-time transport of voice traffic over a Frame Relay network. Service Multiplexing, defined in FRF.11 is a feature that allows for the multiplexing of many voice circuits onto a single VoFR DLCI. The relationship between voice payload sub-frames and Frame Relay frames is illustrated in the following diagram:

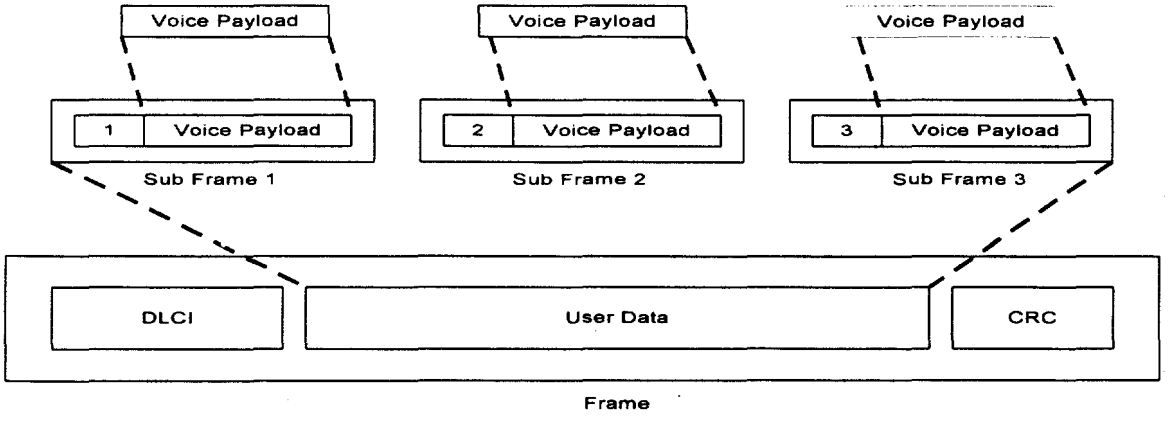

Each sub-frame carries a Sub Channel Identifier that identifies the voice payload sub-frame that is used to carry each individual voice channel.

A Voice Frame Relay Access Device (VFRAD) accomplishes encapsulation of voice traffic into frames. A VFRAD is positioned between a PBX or key set and the Frame Relay network. The VFRAD multiplexes voice, fax, and data from a variety of sources into a common Frame Relay connection. In addition, the VFRAD can provide other services such as compression, encryption, silent suppression, etc. Using voice compression, up to 255 voice sub-channels can be multiplexed within a single Frame Relay DLCI. Addressing is accomplished in VoFR by the transmission of binary representations of dual tone multi-frequency (DTMF). Many aspects of VoFR implementations are left to vendor specific implementations.

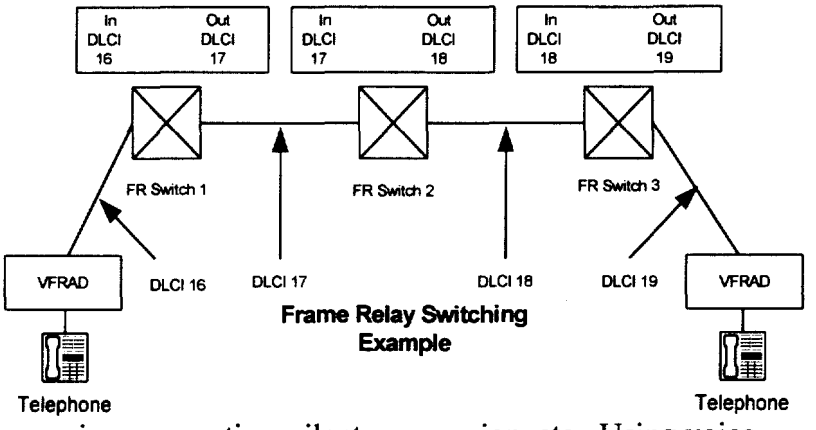

Frame Relay Switching Example

As frames arrive at the ingress to the network, the Frame Relay header is examined to determine the DLCI. The DLCI value is mapped to an outgoing facility, which may use a different DLCI. Note that the DLCI value has only local significance. Each end of the Frame Relay connection can, and probably will utilize a different DLCI value. Nothing about the DLCI value at any node in the network identifies the final destination of the frame and the user data that is being transported.

At no time during the transport of user data does the Frame Relay switch examine the user data to determine anything about the nature of the upper layer protocols, including the nature of voice traffic being transported. The voice sub-channels that are being

carried within the DLCI have significance only at the edges, where they are de-multiplexed by the VFRAD.
It is not technically feasible to examine the user data that is being switched by the Frame Relay switch. Switching nodes are designed to relay the frames as quickly as possible. This design effectively prohibits the Frame Relay switch from examining the user information encapsulated within the frame due to the processing and implementation that would be required. Due to the vendor specific implementation of most aspects of VoFR, having the ability to examine the voice payload could be a moot point because of proprietary compression and encryption.